



Pentium® III Xeon™ Processor Specification Update

Release Date: November 1999

Order Number: 244460-008

The Pentium® III Xeon™ processor may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Pentium® III Xeon™ processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

The Specification Update should be publicly available following the last shipment date for a period of time equal to the specific product's warranty period. Hardcopy Specification Updates will be available for one (1) year following End of Life (EOL). Web access will be available for three (3) years following EOL.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>

Copyright © Intel Corporation 1999.

* Third-party brands and names are the property of their respective owners.

CONTENTS

REVISION HISTORY	ii
PREFACE	iii
Specification Update for the Pentium® III Xeon™ Processor	
GENERAL INFORMATION	1
Pentium III Xeon Processor and Boxed Pentium III Xeon Processor Markings	1
Dynamic Mark Area	1
IDENTIFICATION INFORMATION	2
Mixed Steppings in MP Systems	3
SUMMARY of CHANGES	6
Summary of Errata	7
Summary of Documentation Changes	9
Summary of Specification Clarifications	10
Summary of Specification Changes	10
ERRATA	11
DOCUMENTATION CHANGES	41
SPECIFICATION CLARIFICATIONS	43
SPECIFICATION CHANGES	45



REVISION HISTORY

Date of Revision	Version	Description
March 1999	-001	This document is the first Specification Update for the Pentium® III Xeon™ processor.
April 1999	-002	Removed Erratum G15. Renumbered remaining errata. Added Errata G45 and G46. Corrected Errata table "Plans" column for G29 and G42. Updated the Pentium III Xeon Processor Identification and Package Information table. Moved revised Mixed Steppings statement to the General Information section.
June 1999	-003	Added Errata G47 and G48. . Added Documentation Change G1. Added Specification Clarifications G1 and G2. Added Specification Change G2. Updated the Pentium III Xeon Processor Identification and Package Information table.
July 1999	-004	Added Errata G49 and G50. Added footnote 7 on Pentium III Xeon Processor Identification and Package Information table. Changed the "Plans" column in the Summary Table of Changes from "Fix" to "Fixed" for G35, G36, G37, G44, and G46. Updated Documentation Changes, Specification Clarifications, and Specification Changes introduction paragraphs.
August 1999	-005	Added Errata G51 and G52. Added Documentation Change G2. Updated Codes Used in Summary Table. Updated column heading in Errata, Documentation Changes, Specification Clarifications and Specification Changes tables.
September 1999	-006	Added new Boxed Processors to the Pentium III Xeon Processor Identification and Package table. Revised Erratum G52.
October 1999	-007	Added Brand ID column to <i>Identification Information</i> . Added Errata G53.
November 1999	-008	Revised the <i>Pentium® III Xeon™ Processor Identification Information</i> table. Added errata G54, G55, and G56. Added <i>Specification Change G3</i> . Updated <i>Identification Information</i> , <i>Pentium® III Xeon™ Processor Identification Information</i> , and <i>Mixed Steppings</i> . Added cA2 column to <i>Summary of Errata</i> , <i>Summary of Documentation Changes</i> , <i>Summary of Specification Clarifications</i> , and <i>Summary of Specification Changes</i> .



PREFACE

This document is an update to the specifications contained in the following documents:

- *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet (Order Number 245094)
- *Pentium® III Xeon™ Processor at 600MHz, 667 MHz and 733 MHz* datasheet (Order Number 245094)
- *Intel Architecture Software Developer's Manual, Volumes 1, 2 and 3* (Order Numbers 243190, 243191, and 243192, respectively)

It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools. It contains S-Specs, Errata, Documentation Changes, Specification Clarifications and Specification Changes.

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L2 cache size, package type, etc., as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. Errata may cause the Pentium® III Xeon™ processor's behavior to deviate from published specifications. Hardware and software designed to be used with any given processor must assume that all errata documented for that processor are present on all devices unless otherwise noted.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.


Specification Changes are modifications to the current published specifications for the Pentium III Xeon processor. These changes will be incorporated in the next release of the specifications.

Specification Update for the Pentium® III Xeon™ Processor



GENERAL INFORMATION

Pentium III Xeon Processor and Boxed Pentium III Xeon Processor Markings



Production Dynamic Mark Example:


500/512/100/2.0V S1
FFFFFFFF-NNNN {country}
i®©'98 SSSSS



2D Matrix Contents Example:

Intel 80525PY500512
FFFFFFFF-NNNN

Dynamic Mark Area





IDENTIFICATION INFORMATION

The Pentium III Xeon processor can be identified by the following values:

Family ¹	Model ²	Brand ID ³
0110	0111	00h = Not Supported
0110	1000	03h = "Intel Pentium III Xeon Processor"

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after Reset, bits [11:8] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after Reset, bits [7:4] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
3. The Brand ID corresponds to bits [7:0] of the EBX register after the CUID instruction is executed with a 1 in the EAX register.

The Pentium III Xeon processor's second level (L2) cache size can be determined by the following register contents:

512-Kbyte Unified L2 Cache ¹	43h
1-Mbyte Unified L2 Cache ¹	44h
2-Mbyte Unified L2 Cache ¹	45h
256-Kbyte, 8-way set associative, L2 Cache ¹	82h

NOTE:

1. For the Pentium III Xeon processor, the L2 cache size corresponds to a token in the EDX register after the CUID instruction is executed with a 2 in the EAX register. Other Intel microprocessor models or families may move this information to other bit positions or otherwise reformat the result returned by this instruction; generic code should parse the resulting token stream according to the definition of the CUID instruction.



Mixed Steppings in MP Systems

Intel Corporation fully supports mixed steppings of Pentium III Xeon processors. The following list and processor matrix describes the requirements to support mixed steppings:

- Mixed steppings are only supported with processors that have identical family and model number as indicated by the CPUID instruction.
- While Intel has done nothing to specifically prevent processors operating at differing frequencies from functioning within a multiprocessor system, there may be uncharacterized errata that exist in such configurations. Intel does not support such configurations. In mixed stepping systems, all processors must operate at identical frequencies (i.e., the highest frequency rating commonly supported by all processors).
- While there are no known issues associated with the mixing of processors with differing cache sizes in a multiprocessor system, and Intel has done nothing to specifically prevent such system configurations from operating, Intel does not support such configurations since there may be uncharacterized errata that exist. In mixed stepping systems, all processors must be of the same cache size.
- While Intel believes that certain customers may wish to perform validation of system configurations with mixed frequency or cache sizes, and that those efforts are an acceptable option to our customers, customers would be fully responsible for the validation of such configurations.
- The workarounds identified in this and following specification updates must be properly applied to each processor in the system. Certain errata are specific to the multiprocessor environment and are identified in the *Mixed Stepping Processor Matrix* found at the end of this section. Errata for all processor steppings will affect system performance if not properly worked around. Also see the “Pentium® III Xeon™ Processor Identification and Package Information” section for additional details on which processors are affected by specific errata.
- In mixed stepping systems, the processor with the lowest feature-set, as determined by the CPUID Feature Bytes, must be the Bootstrap Processor (BSP). In the event of a tie in feature-set, the tie should be resolved by selecting the BSP as the processor with the lowest stepping as determined by the CPUID instruction.
- *Functional Redundancy Checking Mode* (FRC mode) is not supported using a master and checker pair of processors with different stepping, model number, cache size, or frequency.

In the following processor matrix, “NI” indicates that there are currently no known issues associated with mixing these steppings. A number indicates that a known issue has been identified as listed in the table following the matrix. A multiprocessor system using mixed processor steppings must assure that errata are addressed appropriately for each processor.

MP/DP Platform Population Matrix for the Pentium III Xeon Processor with 100-MHz System Bus		
Stepping	kB0	kC0
kB0	NI	1, 2
kC0	1, 2	NI

NOTES:

1. Some of these processors are affected by errata, which may affect the features an MP system is able to support. See the “Pentium III Xeon Processor Identification and Package Information” table for details on which processors are affected by these errata.
2. kB0 and kC0 stepping processors may have differing cache components which require differing voltages. When mixing these steppings the proper cache voltage as identified by the VID pins must be supplied to each cartridge.



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

DP Platform Population Matrix for the Pentium III Xeon Processor with 133-MHz System Bus

Pentium III Xeon Processor Stepping	cA2
cA2	NI

Pentium III Xeon Processor Identification and Package Information

S-Spec Number	Core Stepping	CPUID	Speed (MHz)	L2 Size (Kbytes)	Cache and Stepping	Processor Substrate Revision	Cartridge Revision	Notes
SL2XU	kB0	0672h	500/100	512	C6C B0	512K-6A	2.0	1, 2, 5, 7
SL2XV	kB0	0672h	500/100	1024	C6C B0	1M-6A	2.0	1, 2, 5, 7
SL2XW	kB0	0672h	500/100	2048	CK1 B1	2M-kA	2.0	1, 2, 5, 7
SL3C9	kB0	0672h	500/100	512	C6C B0	512K-6A	2.0	1, 2, 3, 5, 7
SL3CA	kB0	0672h	500/100	1024	C6C B0	1M-6A	2.0	1, 2, 3, 5, 7
SL3CB	kB0	0672h	500/100	2048	CK1 B1	2M-kA	2.0	1, 2, 3, 5, 7
SL3FK	kC0	0673h	550/100	512	CK2 B2	512K-kA	2.0	2, 4, 7, 8
SL3D9	kC0	0673h	500/100	512	C6C B0	512K-6A	2.0	2, 7
SL3DA	kC0	0673h	500/100	1024	C6C B0	1M-6A	2.0	2, 7
SL3DB	kC0	0673h	500/100	2048	CK1 B1	2M-kA	2.0	2, 6, 7
SL3AJ	kC0	0673h	550/100	512	CK2 B1	512K-6A	2.0	2, 7, 8
SL3CE	kC0	0673h	550/100	1024	CK1 B1	1M-6A	2.0	2, 7, 8
SL3CF	kC0	0673h	550/100	2048	CK1 B1	2M-kA	2.0	2, 6, 7
SL3TW	kC0	0673h	550/100	1024	CK1 B1	2M-kA	2.0	2, 7
SL3FR	kC0	0673h	550/100	512	CK2 B2	512K-kA	2.0	2, 3, 4, 7, 8
SL385	kC0	0673h	500/100	512	C6C B0	512K-6A	2,0	2, 3, 7
SL386	kC0	0673h	500/100	1024	C6C B0	1M-6A	2,0	2, 3, 7
SL387	kC0	0673h	500/100	2048	CK1 B1	2M-kA	2,0	2, 3, 6, 7
SL3LM	kC0	0673h	550/100	512	CK2 B2	512K-kA	2.0	2, 3, 7, 8
SL3LN	kC0	0673h	550/100	1024	CK1 B1	1M-kA	2.0	2, 3, 7, 8
SL3LP	kC0	0673h	550/100	2048	CK1 B1	2M-kA	2.0	2, 3, 6, 7
SL3BJ	cA2	0681h	600/133	256	N/A	cuB	2.0	9
SL3BK	cA2	0681h	600/133	256	N/A	cuB	2.0	10
SL3BL	cA2	0681h	667/133	256	N/A	cuB	2.0	9
SL3DC	cA2	0681h	667/133	256	N/A	cuB	2.0	10
SL3SF	cA2	0681h	733/133	256	N/A	cuB	2.0	9



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

Pentium III Xeon Processor Identification and Package Information

S-Spec Number	Core Stepping	CPUID	Speed (MHz)	L2 Size (Kbytes)	Cache and Stepping	Processor Substrate Revision	Cartridge Revision	Notes
SL3SG	cA2	0681h	733/133	256	N/A	cuB	2.0	10
SL3SS	cA2	0681h	600/133	256	N/A	cuB	2.0	3, 9
SL3ST	cA2	0681h	667/133	256	N/A	cuB	2.0	3, 9
SL3SU	cA2	0681h	733/133	256	N/A	cuB	2.0	3, 9

NOTES:

1.

These processors are affected by Erratum G44.
2.

The performance-monitoring event counter 1 may be inaccurate when counting events for the Data Cache Unit and External Bus Logic in these processors. Use counter 0 to count these events.
3.

This is a boxed processor with attached passive heatsink.
4.

These processors are validated for use in two-way systems only.
5.

These processors are affected by Erratum G46.
6.

These processors are affected by Erratum G48.
7.

Performance-monitoring event counters do not reflect MOVD and MOVQ stores to memory on these processors.
8.

These processors are affected by Erratum G56.
9.

These processors are designed to operate at 2.8V, and have been tested to 55°C T_{PLATE}.
10.

These processors are designed to operate at either 5V or 12V, and have been tested to 55°C T_{PLATE}.



SUMMARY of CHANGES

The following table indicates the Errata, Documentation Changes, Specification Clarifications, or Specification Changes that apply to Pentium III Xeon processors. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Erratum, Documentation Change, Specification Clarification, or Specification Change applies to the given processor stepping.
(No mark) or (blank box):	This item is fixed in or does not apply to the given stepping.
Fix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
Doc:	Intel intends to update the appropriate documentation in a future revision.
PKG:	This column refers to errata on the Pentium III Xeon processor substrate.
AP:	APIC related erratum.
Shaded:	This item is either new or modified from the previous version of the document.

Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

- A = Pentium® II processor
- B = Mobile Pentium® II processor
- C = Intel® Celeron™ processor
- D = Pentium® II Xeon™ processor
- E = Pentium® III processor
- G = Pentium® III Xeon™ processor
- H = Intel® Mobile Celeron™ processor

The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

Summary of Errata

NO.	kB0	kc0	cA2	PKG	Plans	ERRATA
G1	X	X	X		NoFix	FP Data Operand Pointer may be incorrectly calculated after FP access which wraps 64-Kbyte boundary in 16-bit code
G2	X	X	X		NoFix	Differences exist in debug exception reporting
G3	X	X	X		NoFix	FLUSH# servicing delayed while waiting for STARTUP_IPI in MP systems
G4	X	X	X		NoFix	Code fetch matching disabled debug register may cause debug exception
G5	X	X	X		NoFix	Double ECC error on read may result in INIT#
G6	X	X	X		NoFix	FP inexact-result exception flag may not be set
G7	X	X	X		NoFix	BTM for SMI will contain incorrect FROM EIP
G8	X	X	X		NoFix	I/O restart in SMM may fail after simultaneous MCE
G9	X	X	X		NoFix	Branch traps do not function if BTMs are also enabled
G10	X	X			NoFix	Checker BIST failure in FRC mode not signaled
G11	X	X			NoFix	INIT# assertion causes FRCERR assertion in FRC mode
G12	X	X	X		NoFix	Machine check exception handler may not always execute successfully
G13	X	X	X		NoFix	LBERR may be corrupted after some events
G14	X	X	X		NoFix	BTMs may be corrupted during simultaneous L1 cache line replacement
G15	X	X	X		NoFix	Near CALL to ESP creates unexpected EIP address
G16	X	X	X		NoFix	Mixed cacheability of lock variables problematic in MP systems
G17	X	X	X		NoFix	MCE due to L2 parity error gives L1 MCACOD.LL
G18	X	X	X		NoFix	Memory Type field undefined for nonmemory operations
G19	X	X	X		NoFix	Infinite snoop stall during L2 initialization of MP systems
G20	X	X	X		NoFix	FP Data Operand Pointer may not be zero after power on or Reset
G21	X	X	X		NoFix	Premature execution of a load operation prior to exception handler invocation
G22	X	X	X		NoFix	EFLAGS discrepancy on page fault after multiprocessor TLB shutdown

**PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE**

NO.	kB0	kc0	cA2	PKG	Plans	ERRATA
G23	X	X	X		NoFix	Read portion of RMW instruction may execute twice
G24	X	X	X		NoFix	MC2_STATUS MSR has model-specific error code and Machine Check Architecture error code reversed
G25	X	X	X		NoFix	MOVD following zeroing instruction can cause incorrect result
G26	X	X	X		NoFix	Top 4 PAT entries not usable with Mode B or Mode C paging
G27	X	X	X		NoFix	MOV with debug register causes debug exception
G28	X	X	X		NoFix	Data Breakpoint Exception in a displacement relative near call may corrupt EIP
G29	X	X	X		NoFix	System bus ECC not functional with 2:1 ratio
G30	X	X	X		NoFix	RDMSR or WRMSR to invalid MSR address may not cause GP fault
G31	X	X	X		NoFix	SYSENTER/SYSEXIT instructions can implicitly load "null segment selector" to SS and CS registers
G32	X	X	X		NoFix	PRELOAD followed by EXTEST does not load boundary scan data
G33	X	X	X		NoFix	Far jump to new TSS with D-bit cleared may cause system hang
G34	X	X	X		NoFix	Illegal opcode during L2 cache initialization
G35	X				Fixed	Incorrect L2 cache line invalidation
G36	X				Fixed	Transmission error on cache read
G37	X				Fixed	COMISS/UCOMISS may not update EFLAGS under certain conditions
G38	X	X	X		NoFix	System hang may occur with 2:1 core to bus ratio
G39	X	X	X		NoFix	Misaligned locked access to APIC space results in hang
G40	X	X	X		NoFix	Potential loss of data coherency during MP data ownership transfer
G41	X	X	X		NoFix	INT 1 instruction handler execution could generate a debug exception
G42	X	X	X		NoFix	Memory ordering-based synchronization may cause a livelock condition in MP systems
G43	X	X	X		NoFix	Floating point exception signal may be deferred
G44	X				Fixed	System bus address parity checking may report false AERR#
G45	X	X	X		NoFix	Processor may assert DRDY# on a write with no data



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

NO.	kB0	kc0	cA2	PKG	Plans	ERRATA
G46				X	Fixed	Thermal sensor leakage current may exceed specification
G47	X	X	X		NoFix	GP# fault on WRMSR to ROB_CR_BKUPTMPDR6
G48				X	NoFix	Heavy L2 cache traffic results in noise on the TCK signal
G49	X	X	X		NoFix	Machine check exception may occur due to improper line eviction in the IFU
G50	X	X	X		NoFix	Machine check exception may occur during L2 cache initialization.
G51	X	X	X		NoFix	Snoop request may cause DBSY# hang
G52	X	X	X		Fix	Performance counters include Streaming SIMD Extensions L1 prefetch
G53	X	X	X		NoFix	Lower bits of SMRAM SMBASE register cannot be written with an ITP
G54	X	X	X		Fix	Task switch caused by page fault may cause wrong PTE and PDE access bit to be set
G55	X	X	X		NoFix	Cross-modifying code operations on a jump instruction may cause general protection fault
G56				X	NoFix	V _{IH} specification deviation on processor RS0# input
G1AP	X	X	X		NoFix	APIC access to cacheable memory causes shutdown
G2AP	X	X	X		NoFix	MP systems may hang due to catastrophic errors during BSP determination
G3AP	X	X	X		NoFix	Write to mask LVT (programmed as EXTINT) will not deassert outstanding interrupt

NOTE: The erratum titled *A20M# may be inverted after returning from SMM and Reset* and numbered as G15 has been removed. It was erroneously included in previous versions of this Specification Update. The remaining errata have been renumbered.

Summary of Documentation Changes

NO.	KB0	kC0	cA2	PKG	Plans	DOCUMENTATION CHANGES
G1	X	X	X		Doc	STPCLK# pin definition
G2	X	X	X		Doc	Invalidating caches and TLBs
G3	X	X	X		Doc	Handling of self-modifying and cross-modifying code



Summary of Specification Clarifications

NO.	kB0	kC0	cA2	PKG	Plans	SPECIFICATION CLARIFICATIONS
G1	X	X	X		Doc	MTRR initialization clarification
G2	X	X	X		Doc	Floating point opcode clarification

Summary of Specification Changes

NO.	kB0	kC0	cA2	PKG	Plans	SPECIFICATION CHANGES
G1	X	X	X		Doc	Locks across cache line boundary disable bit added
G2	X	X	X		Doc	RESET# pin definition



ERRATA

G1. *FP Data Operand Pointer May Be Incorrectly Calculated After FP Access Which Wraps 64-Kbyte Boundary in 16-Bit Code*

Problem: The FP Data Operand Pointer is the effective address of the operand associated with the last noncontrol floating-point instruction executed by the machine. If an 80-bit floating-point access (load or store) occurs in a 16-bit mode other than protected mode (in which case the access will produce a segment limit violation), the memory access wraps a 64-Kbyte boundary, and the floating-point environment is subsequently saved, the value contained in the FP Data Operand Pointer may be incorrect.

Implication: A 32-bit operating system running 16-bit floating-point code may encounter this erratum, under the following conditions:

- The operating system is using a segment greater than 64 Kbytes in size.
- An application is running in a 16-bit mode other than protected mode.
- An 80-bit floating-point load or store which wraps the 64-Kbyte boundary is executed.
- The operating system performs a floating-point environment store (FSAVE/FNSAVE/FSTENV/FNSTENV) after the above memory access.
- The operating system uses the value contained in the FP Data Operand Pointer.

Wrapping an 80-bit floating-point load around a segment boundary in this way is not a normal programming practice. Intel has not currently identified any software which exhibits this behavior.

Workaround: If the FP Data Operand Pointer is used in an OS which may run 16-bit floating-point code, care must be taken to ensure that no 80-bit floating-point accesses are wrapped around a 64-Kbyte boundary.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G2. *Differences Exist in Debug Exception Reporting*

Problem: There exist some differences in the reporting of code and data breakpoint matches between that specified by previous Intel processors' specifications and the behavior of the Pentium III Xeon processor, as described below:

Case 1: The first case is for a breakpoint set on a MOVSS or POPSS instruction, when the instruction following it causes a debug register protection fault (DR7.gd is already set, enabling the fault). The Pentium processor reports delayed data breakpoint matches from the MOVSS or POPSS instructions by setting the matching DR6.bi bits, along with the debug register protection fault (DR6.bd). If additional breakpoint faults are matched during the call of the debug fault handler, the Pentium processor sets the breakpoint match bits (DR6.bi) to reflect the breakpoints matched by both the MOVSS or POPSS breakpoint and the debug fault handler call. The Pentium III Xeon processor only sets DR6.bd in either situation, and does not set any of the DR6.bi bits.

Case 2: In the second breakpoint reporting failure case, if a MOVSS or POPSS instruction with a data breakpoint is followed by a store to memory which crosses a 4-Kbyte page boundary, the breakpoint information for the MOVSS or POPSS will be lost. Previous processors retain this information across such a page split.



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

Case 3: If they occur after a MOVSS or POPSS instruction, the INT *n*, INTO, and INT3 instructions zero the DR6.Bi bits (bits B0 through B3), clearing pending breakpoint information, unlike previous processors.

Case 4: If a data breakpoint and an SMI (System Management Interrupt) occur simultaneously, the SMI will be serviced via a call to the SMM handler, and the pending breakpoint will be lost.

Case 5: When an instruction which accesses a debug register is executed, and a breakpoint is encountered on the instruction, the breakpoint is reported twice.

Implication: When debugging or when developing debuggers for a Pentium III Xeon processor-based system, this behavior should be noted. Normal usage of the MOVSS or POPSS instructions (i.e., following them with a MOV ESP) will not exhibit the behavior of cases 1-3. Debugging in conjunction with SMM will be limited by case 4.

Workaround: Following MOVSS and POPSS instructions with a MOV ESP instruction when using breakpoints will avoid the first three cases of this erratum. No workaround has been identified for cases 4 or 5.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G3. FLUSH# Servicing Delayed While Waiting for STARTUP_IPI in MP Systems

Problem: In an MP system, if an application processor is waiting for a startup inter-processor interrupt (STARTUP_IPI), then it will not service a FLUSH# pin assertion until it has received the STARTUP_IPI.

Implication: After the MP initialization protocol, only one processor becomes the bootstrap processor (BSP). The other processor becomes a slave application processor (AP). After losing the BSP arbitration, the AP goes into a wait loop, waiting for a STARTUP_IPI.

The BSP can wake up the AP to perform some tasks with a STARTUP_IPI, and then put it back to sleep with an initialization inter-processor interrupt (INIT_IPI, which has the same effect as asserting INIT#), which returns it to a wait loop. The result is a possible loss of cache coherency if the off-line processor is intended to service a FLUSH# assertion at this point. The FLUSH# will be serviced as soon as the processor is awakened by a STARTUP_IPI, before any other instructions are executed. Intel has not encountered any operating systems that are affected by this erratum.

Workaround: Operating system developers should take care to execute a WBINVD instruction before the AP is taken off-line using an INIT_IPI.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G4. Code Fetch Matching Disabled Debug Register May Cause Debug Exception

Problem: The bits L0-3 and G0-3 enable breakpoints local to a task and global to all tasks, respectively. If one of these bits is set, a breakpoint is enabled, corresponding to the addresses in the debug registers DR0-DR3. If at least one of these breakpoints is enabled, any of these registers are *disabled* (i.e., L_n and G_n are 0), and RW_n for the disabled register is 00 (indicating a breakpoint on instruction execution), normally an instruction fetch will not cause an instruction-breakpoint fault based on a match with the address in the disabled register(s). However, if the address in a disabled register matches the address of a code fetch which also results in a page fault, an instruction-breakpoint fault will occur.

Implication: While debugging software, extraneous instruction-breakpoint faults may be encountered if breakpoint registers are not cleared when they are disabled. Debug software which does not implement a code breakpoint handler will fail, if this occurs. If a handler is present, the fault will be serviced. Mixing data and code may exacerbate this problem by allowing disabled data breakpoint registers to break on an instruction fetch.

Workaround: The debug handler should clear breakpoint registers before they become disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G5. Double ECC Error on Read May Result in BINIT#

Problem: For this erratum to occur, the following conditions must be met:

- Machine Check Exceptions (MCEs) must be enabled.
- A dataless transaction (such as a write invalidate) must be occurring simultaneously with a transaction which returns data (a normal read).
- The read data must contain a double-bit uncorrectable ECC error.

If these conditions are met, the Pentium III Xeon processor will not be able to determine which transaction was erroneous, and instead of generating an MCE, it will generate a BINIT#.

Implication: The bus will be reinitialized in this case. However, since a double-bit uncorrectable ECC error occurred on the read, the MCE handler (which is normally reached on a double-bit uncorrectable ECC error for a read) would most likely cause the same BINIT# event.

Workaround: Though the ability to drive BINIT# can be disabled in the Pentium III Xeon processor, which would prevent the effects of this erratum, overall system behavior would not improve, since the error which would normally cause a BINIT# would instead cause the machine to shut down. No other workaround has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G6. *FP Inexact-Result Exception Flag May Not Be Set*

Problem: When the result of a floating-point operation is not exactly representable in the destination format (1/3 in binary form, for example), an inexact-result (precision) exception occurs. When this occurs, the PE bit (bit 5 of the FPU status word) is normally set by the processor. Under certain rare conditions, this bit may not be set when this rounding occurs. However, other actions taken by the processor (invoking the software exception handler if the exception is unmasked) are not affected. This erratum can only occur if the floating-point operation which causes the precision exception is immediately followed by one of the following instructions:

- FST m32real
- FST m64real
- FSTP m32real
- FSTP m64real
- FSTP m80real
- FIST m16int
- FIST m32int
- FISTP m16int
- FISTP m32int
- FISTP m64int

Note that even if this combination of instructions is encountered, there is also a dependency on the internal pipelining and execution state of both instructions in the processor.

Implication: Inexact-result exceptions are commonly masked or ignored by applications, as it happens frequently, and produces a rounded result acceptable to most applications. The PE bit of the FPU status word may not always be set upon receiving an inexact-result exception. Thus, if these exceptions are unmasked, a floating-point error exception handler may not recognize that a precision exception occurred. Note that this is a “sticky” bit, i.e., once set by an inexact-result condition, it remains set until cleared by software.

Workaround: This condition can be avoided by inserting two NOP instructions between the two floating-point instructions.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G7. *BTM for SMI Will Contain Incorrect FROM EIP*

Problem: A system management interrupt (SMI) will produce a Branch Trace Message (BTM), if BTMs are enabled. However, the FROM EIP field of the BTM (used to determine the address of the instruction which was being executed when the SMI was serviced) will not have been updated for the SMI, so the field will report the same FROM EIP as the previous BTM.

Implication: A BTM which is issued for an SMI will not contain the correct FROM EIP, limiting the usefulness of BTMs for debugging software in conjunction with System Management Mode (SMM).

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G8. I/O Restart in SMM May Fail After Simultaneous MCE

Problem: If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the Pentium III Xeon processor will signal a machine check exception (MCE). If the instruction is directed at a device which is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, upon attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is completed successfully, it will attempt to restart the I/O instruction, but will not have the correct machine state, due to the call to the MCE handler.

Implication: A simultaneous MCE and SMI# assertion may occur for one of the I/O instructions above. The SMM handler may attempt to restart such an I/O instruction, but will have corrupted state due to the MCE handler call, leading to failure of the restart and shutdown of the processor.

Workaround: If a system implementation must support both SMM and MCEs, the first thing the SMM handler code (when an I/O restart is to be performed) should do is check for a pending MCE. If there is an MCE pending, the SMM handler should immediately exit via an RSM instruction and allow the machine check exception handler to execute. If there is not, the SMM handler may proceed with its normal operation.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G9. Branch Traps Do Not Function if BTMs Are Also Enabled

Problem: If branch traps or branch trace messages (BTMs) are enabled alone, both function as expected. However, if both are enabled, only the BTMs will function, and the branch traps will be ignored.

Implication: The branch traps and branch trace message debugging features cannot be used together.

Workaround: If branch trap functionality is desired, BTMs must be disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G10. Checker BIST Failure in FRC Mode Not Signaled

Problem: If a system is running in functional redundancy checking (FRC) mode, and the checker of the master-checker pair encounters a hard failure while running the built-in self test (BIST), the checker will tri-state all outputs without signaling an IERR#.

Implication: Assuming the master passes BIST successfully, it will continue execution unchecked, operating without functional redundancy. However, the necessary pull-up on the FRCERR pin will cause an FRCERR to be signaled. The operation of the master depends on the implementation of FRCERR.

Workaround: For successful detection of BIST failure in the checker of an FRC pair, use the FRCERR signal, instead of IERR#.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G11. *BINIT# Assertion Causes FRCERR Assertion in FRC Mode*

Problem: If a pair of Pentium III Xeon processors are running in functional redundancy checking (FRC) mode, and a catastrophic error condition causes BINIT# to be asserted, the checker in the master-checker pair will enter shutdown. The next bus transaction from the master will then result in the assertion of FRCERR.

Implication: Bus initialization via an assertion of BINIT# occurs as the result of a catastrophic error condition which precludes the continuing reliable execution of the system. Under normal circumstances, the master-checker pair would remain synchronized in the execution of the BINIT# handler. However, due to this erratum, an FRCERR will be signaled. System behavior then depends on the system specific error recovery mechanisms.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G12. *Machine Check Exception Handler May Not Always Execute Successfully*

Problem: An asynchronous machine check exception (MCE), such as a BINIT# event, which occurs during an access that splits a 4-Kbyte page boundary may leave some internal registers in an indeterminate state. Thus, MCE handler code may not always run successfully if an asynchronous MCE has occurred previously.

Implication: An MCE may not always result in the successful execution of the MCE handler. However, asynchronous MCEs usually occur upon detection of a catastrophic system condition that would also hang the processor. Leaving MCEs disabled will result in the condition which caused the asynchronous MCE instead causing the processor to enter shutdown. Therefore, leaving MCEs disabled may not improve overall system behavior.

Workaround: No workaround which would guarantee successful MCE handler execution under this condition has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G13. ***LBER May Be Corrupted After Some Events***

Problem: The last branch record (LBR) and the last branch before exception record (LBER) can be used to determine the source and destination information for previous branches or exceptions. The LBR contains the source and destination addresses for the last branch or exception, and the LBER contains similar information for the last branch taken before the last exception. This information is typically used to determine the location of a branch which leads to execution of code which causes an exception. However, after a catastrophic bus condition which results in an assertion of BINIT# and the re-initialization of the buses, the value in the LBER may be corrupted. Also, after either a CALL which results in a fault or a software interrupt, the LBER and LBR will be updated to the same value, when the LBER should not have been updated.

Implication: The LBER and LBR registers are used only for debugging purposes. When this erratum occurs, the LBER will not contain reliable address information. The value of LBER should be used with caution when debugging branching code; if the values in the LBR and LBER are the same, then the LBER value is incorrect. Also, the value in the LBER should not be relied upon after a BINIT# event.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G14. ***BTMs May Be Corrupted During Simultaneous L1 Cache Line Replacement***

Problem: When Branch Trace Messages (BTMs) are enabled and such a message is generated, the BTM may be corrupted when issued to the bus by the L1 cache if a new line of data is brought into the L1 data cache simultaneously. Though the new line being stored in the L1 cache is stored correctly, and no corruption occurs in the data, the information in the BTM may be incorrect due to the internal collision of the data line and the BTM.

Implication: Although BTMs may not be entirely reliable due to this erratum, the conditions necessary for this boundary condition to occur have only been exhibited during focused simulation testing. Intel has currently not observed this erratum in a system level validation environment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G15. Near CALL to ESP Creates Unexpected EIP Address

Problem: As documented, the CALL instruction saves procedure linking information in the procedure stack and jumps to the called procedure specified with the destination (target) operand. The target operand specifies the address of the first instruction in the called procedure. This operand can be an immediate value, a general purpose register, or a memory location. When accessing an absolute address indirectly using the stack pointer (ESP) as a base register, the base value used is the value in the ESP register before the instruction executes. However, when accessing an absolute address directly using ESP as the base register, the base value used is the value of ESP *after* the return value is pushed on the stack, not the value in the ESP register *before* the instruction executed.

Implication: Due to this erratum, the processor may transfer control to an unintended address. Results are unpredictable, depending on the particular application, and can range from no effect to the unexpected termination of the application due to an exception. Intel has observed this erratum only in a focused testing environment. Intel has not observed any commercially available operating system, application, or compiler that makes use of or generates this instruction.

Workaround: If the other seven general purpose registers are unavailable for use, and it is necessary to do a CALL via the ESP register, first push ESP onto the stack, then perform an *indirect* call using ESP (e.g., CALL [ESP]). The saved version of ESP should then be popped off the stack after the call returns.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G16. Mixed Cacheability of Lock Variables Is Problematic in MP Systems

Problem: This errata only affects multiprocessor systems where a lock variable address is marked cacheable in one processor and uncacheable in any others. The processors which have it marked uncacheable may stall indefinitely when accessing the lock variable. The stall is only encountered if:

- One processor has the lock variable cached, and is attempting to execute a cache lock.
- If the processor which has that address cached has it cached in its L2 only.
- Other processors, meanwhile, issue back to back accesses to that same address on the bus.

Implication: MP systems where all processors either use cache locks or consistent locks to uncacheable space will not encounter this problem. If, however, a lock variable's cacheability varies in different processors, and several processors are all attempting to perform the lock simultaneously, an indefinite stall may be experienced by the processors which have it marked uncacheable in locking the variable (if the conditions above are satisfied). Intel has only encountered this problem in focus testing with artificially generated external events. Intel has not currently identified any commercial software which exhibits this problem.

Workaround: Follow a homogenous model for the memory type range registers (MTRRs), ensuring that all processors have the same cacheability attributes for each region of memory; do not use locks whose memory type is cacheable on one processor, and uncacheable on others. Avoid page table aliasing, which may produce a nonhomogenous memory model.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G17. MCE Due to L2 Parity Error Gives L1 MCACOD.LL

Problem: If a Cache Reply Parity (CRP) error, Cache Address Parity (CAP) error, or Cache Synchronous Error (CSER) occurs on an access to the Pentium III Xeon processor's L2 cache, the resulting Machine Check Architectural Error Code (MCACOD) will be logged with '01' in the LL field. This value indicates an L1 cache error; the value should be '10', indicating an L2 cache error. Note that L2 ECC errors have the correct value of '10' logged.

Implication: An L2 cache access error, other than an ECC error, will be improperly logged as an L1 cache error in MCACOD.LL.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G18. Memory Type Field Undefined for Nonmemory Operations

Problem: The Memory Type field for nonmemory transactions such as I/O and Special Cycles are undefined. Although the Memory Type attribute for nonmemory operations logically should (and usually does) manifest itself as UC, this feature is not designed into the implementation and is therefore inconsistent.

Implication: Bus agents may decode a non-UC memory type for nonmemory bus transactions.

Workaround: Bus agents must consider transaction type to determine the validity of the Memory Type field for a transaction.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G19. Infinite Snoop Stall During L2 Initialization of MP Systems

Problem: It is possible for snoop traffic generated on the system bus while a processor is executing its L2 cache initialization routine to cause the initializing processor to hang.

Implication: An MP system which does not suppress snoop traffic while L2 caches are being initialized may hang during this initialization sequence.

Workaround: System BIOS can create an execution environment which allows processors to initialize their L2 caches without the system generating any snoop traffic on the bus.

Below is a pseudo-code fragment, designed explicitly for a four-processor system, that uses a serial algorithm to initialize each processor's L2 cache:

```
Suppress_all_I/O_traffic()
K = 0;
while (K <= 3)
{
    /* Obtain current value of K. This forces both Temp and K into */
    /* the L1 cache. Note that Temp could also be maintained in a */
    /* general purpose register. */

    Temp = K;
    Wait_until_all_processors_are_signed_in_at_barrier()
```

```

if ( logical_proc_APIC_id == K ) {
{
wait_10_usecs_delay_loop(); /* this time delay, required */
/* in the worst case, allows */
/* the barrier semaphore to */
/* settle to shared state. */
Initialize L2 cache
K++
}
else
while (Temp == K);
}
}

```

This algorithm prevents bus snoop traffic from the other processors, which would otherwise cause the initializing processor to hang. The algorithm assumes that the L1 cache is enabled (the Temp and K variables must be cached by each processor). Also, the Memory Type Range Register (MTRR) for the data segment must be set to WB (writeback) memory type.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G20. FP Data Operand Pointer May Not Be Zero After Power On or Reset

Problem: The FP Data Operand Pointer, as specified, should be reset to zero upon power on or Reset by the processor. Due to this erratum, the FP Data Operand Pointer may be nonzero after power on or Reset.

Implication: Software which uses the FP Data Operand Pointer and count on its value being zero after power on or Reset without first executing an FINIT/FNINIT instruction will use an incorrect value, resulting in incorrect behavior of the software.

Workaround: Software should follow the recommendation in Section 8.2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide* (Order Number 243192). This recommendation states that if the FPU will be used, software-initialization code should execute an FINIT/FNINIT instruction following a hardware reset. This will correctly clear the FP Data Operand Pointer to zero.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G21. Premature Execution of Load Operation Prior to Exception Handler Invocation

Problem: This erratum can occur in any of the following situations:

1. If an instruction that performs a memory load causes a code segment limit violation,
2. If a waiting floating-point instruction or MMX™ instruction that performs a memory load has a floating-point exception pending, or
3. If an MMX instruction that performs a memory load and has either CR0.EM = 1 (Emulation bit set), a floating-point Top-of-Stack (FP TOS) not equal to 0, or a DNA exception pending.



If any of the above circumstances occur, it is possible that the load portion of the instruction will have executed before the exception handler is entered.

Implication: In normal code execution where the target of the load operation is to write back memory there is no impact from the load being prematurely executed, nor from the restart and subsequent re-execution of that instruction by the exception handler. If the target of the load is to uncached memory that has a system side-effect, restarting the instruction may cause unexpected system behavior due to the repetition of the side-effect.

Workaround: Code which performs loads from memory that has side-effects can effectively work around this behavior by using simple integer-based load instructions when accessing side-effect memory, and by ensuring that all code is written such that a code segment limit violation cannot occur as a part of reading from side-effect memory.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G22. EFLAGS Discrepancy on Page Fault After Multiprocessor TLB Shutdown

Problem: This erratum may occur when the Pentium III Xeon processor executes one of the following read-modify-write arithmetic instructions and a page fault occurs during the store of the memory operand: ADD, AND, BTC, BTR, BTS, CMPXCHG, DEC, INC, NEG, NOT, OR, ROL/ROR, SAL/SAR/SHL/SHR, SHLD, SHRD, SUB, XOR, and XADD. In this case, the EFLAGS value pushed onto the stack of the page fault handler may reflect the status of the register after the instruction would have completed execution rather than before it. The following conditions are required for the store to generate a page fault and call the operating system page fault handler:

1. The store address entry must be evicted from the DTLB by speculative loads from other instructions that hit the same way of the DTLB before the store has completed. DTLB eviction requires at least three load operations that have linear address bits 15:12 equal to each other and address bits 31:16 different from each other in close physical proximity to the arithmetic operation.
2. The page table entry for the store address must have its permissions tightened during the very small window of time between the DTLB eviction and execution of the store. Examples of page permission tightening include from Present to Not Present or from Read/Write to Read Only, etc.
3. Another processor, without corresponding synchronization and TLB flush, must cause the permission change.

Implication: This scenario may only occur on a multiprocessor platform running an operating system that performs "lazy" TLB shutdowns. The memory image of the EFLAGS register on the page fault handler's stack prematurely contains the final arithmetic flag values although the instruction has not yet completed. Intel has not identified any operating systems that inspect the arithmetic portion of the EFLAGS register during a page fault nor observed this erratum in laboratory testing of software applications.

Workaround: No workaround is needed upon normal restart of the instruction, since this erratum is transparent to the faulting code and results in correct instruction behavior. Operating systems may ensure that no processor is currently accessing a page that is scheduled to have its page permissions tightened or have a page fault handler that ignores any incorrect state.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G23. *Read Portion of RMW Instruction May Execute Twice*

Problem: When the Pentium III Xeon processor executes a read-modify-write (RMW) arithmetic instruction, with memory as the destination, it is possible for a page fault to occur during the execution of the store on the memory operand after the read operation has completed but before the write operation completes. If the memory targeted for the instruction is UC (uncached), memory will observe the occurrence of the initial load before the page fault handler and again if the instruction is restarted.

Implication: If the memory targeted for the RMW instruction has no side effects, then the memory location will simply be read twice with no additional implications. If, however, the load targets a memory region that has side effects, multiple occurrences of the initial load may lead to unpredictable system behavior.

Workaround: Hardware and software developers who write device drivers for custom hardware that may have a side effect style of design should use simple loads and simple stores to transfer data to and from the device.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G24. *MC2_STATUS MSR Has Model-Specific Error Code and Machine Check Architecture Error Code Reversed*

Problem: The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, documents that for the MC_i_STATUS MSR, bits 15:0 contain the MCA (machine-check architecture) error code field, and bits 31:16 contain the model-specific error code field. However, for the MC2_STATUS MSR, these bits have been reversed. For the MC2_STATUS MSR, bits 15:0 contain the model-specific error code field and bits 31:16 contain the MCA error code field.

Implication: A machine check error may be decoded incorrectly if this erratum on the MC2_STATUS MSR is not taken into account.

Workaround: When decoding the MC2_STATUS MSR, reverse the two error fields.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G25. *MOVD Following Zeroing Instruction Can Cause Incorrect Result*

Problem: An incorrect result may be calculated after the following circumstances occur:

1. A register has been zeroed with either a SUB reg, reg instruction or an XOR reg, reg instruction,
2. A value is moved with sign extension into the same register's lower 16 bits; or a signed integer multiply is performed to the same register's lower 16 bits,
3. This register is then copied to an MMX™ technology register using the MOVD instruction prior to any other operations on the sign-extended value.

Specifically, the sign may be incorrectly extended into bits 16-31 of the MMX technology register. Only the MMX technology register is affected by this erratum.

The erratum only occurs when the 3 following steps occur in the order shown. The erratum may occur with up to 40 intervening instructions that do not modify the sign-extended value between steps 2 and 3.

1. XOR EAX, EAX
or SUB EAX, EAX
2. MOVSBX AX, BL
or MOVSBX AX, byte ptr <memory address> or MOVSBX AX, BX
or MOVSBX AX, word ptr <memory address> or IMUL BL (AX implicit, opcode F6 /5)
or IMUL byte ptr <memory address> (AX implicit, opcode F6 /5) or IMUL AX, BX (opcode 0F AF /r)
or IMUL AX, word ptr <memory address> (opcode 0F AF /r) or IMUL AX, BX, 16 (opcode 6B /r ib)
or IMUL AX, word ptr <memory address>, 16 (opcode 6B /r ib) or IMUL AX, 8 (opcode 6B /r ib)
or IMUL AX, BX, 1024 (opcode 69 /r iw)
or IMUL AX, word ptr <memory address>, 1024 (opcode 69 /r iw) or IMUL AX, 1024 (opcode 69 /r iw)
or CBW
3. MOVD MM0, EAX

Note that the values for immediate byte/words are merely representative (i.e., 8, 16, 1024) and that any value in the range for the size is affected. Also, note that this erratum may occur with "EAX" replaced with any 32-bit general purpose register, and "AX" with the corresponding 16-bit version of that replacement. "BL" or "BX" can be replaced with any 8-bit or 16-bit general purpose register. The CBW and IMUL (opcode F6 /5) instructions are specific to the EAX register only.

In the example, EAX is forced to contain 0 by the XOR or SUB instructions. Since the four types of the MOVSBX or IMUL instructions and the CBW instruction modify only bits 15:8 of EAX by sign extending the lower eight bits of EAX, bits 31:16 of EAX should always contain 0. This implies that when MOVD copies EAX to MM0, bits 31:16 of MM0 should also be 0. Under certain scenarios, bits 31:16 of MM0 are not 0, but are replicas of bit 15 (the 16th bit) of AX. This is noticeable when the value in AX after the MOVSBX, IMUL or CBW instruction is negative, i.e., bit 15 of AX is a 1.

When AX is positive (bit 15 of AX is a 0), MOVD will always produce the correct answer. If AX is negative (bit 15 of AX is a 1), MOVD may produce the right answer or the wrong answer depending on the point in time when the MOVD instruction is executed in relation to the MOVSBX, IMUL or CBW instruction.

Implication: The effect of incorrect execution will vary from unnoticeable, due to the code sequence discarding the incorrect bits, to an application failure. If the MMX technology-enabled application in which MOVD is used to manipulate pixels, it is possible for one or more pixels to exhibit the wrong color or position momentarily. It is also possible for a computational application that uses the MOVD instruction in the manner described above to produce incorrect data. Note that this data may cause an unexpected page fault or general protection fault.



Workaround: There are two possible workarounds for this erratum:

1. Rather than using the MOVXSX-MOVD, IMUL-MOVD or CBW-MOVD pairing to handle one variable at a time, use the sign extension capabilities (PSRAW, etc.) within MMX technology for operating on multiple variables. This would result in higher performance as well.
2. Insert another operation that modifies or copies the sign-extended value between the MOVXSX/IMUL/CBW instruction and the MOVD instruction as in the example below:
XOR EAX, EAX (or SUB EAX, EAX)
MOVXSX AX, BL (or other MOVXSX, other IMUL or CBW instruction)
*MOV EAX, EAX
MOVD MM0, EAX

*Note: MOV EAX, EAX is used here as it is fairly generic. Again, EAX can be any 32-bit register.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G26. Top 4 PAT Entries Not Usable With Mode B or Mode C Paging

Problem: The Page Attribute Table (PAT) contains eight entries, which must all be initialized and considered when setting up memory types for the Pentium III Xeon processor. However, in Mode B or Mode C paging, the top four entries do not function correctly for 4-Kbyte pages. Specifically, bit 7 of page table entries which translate addresses to 4-Kbyte pages should be used as the upper bit of a three-bit index to determine the PAT entry that specifies the memory type for the page. When Mode B (CR4.PSE = 1) and/or Mode C (CR4.PAE = 1) are enabled, the processor forces this bit to zero when determining the memory type, regardless of the value in the page table entry. The upper four entries of the PAT function correctly for 2-Mbyte and 4-Mbyte large pages (specified by bit 12 of the page directory entry for those translations).

Implication: Only the lower four PAT entries are useful for 4-Kbyte translations when Mode B or C paging is used. In Mode A paging (4-Kbyte pages only), all eight entries may be used. All eight entries may also be used for large pages in Mode B or C paging.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G27. *MOV with Debug Register Causes Debug Exception*

Problem: When in V86 mode, if a MOV instruction is executed on debug registers, a general-protection exception (#GP) should be generated, as documented in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, Section 14.2. However, in the case when the general detect enable flag (GD) bit is set, the observed behavior is that a debug exception (#DB) is generated instead.

Implication: With debug-register protection enabled (i.e., the GD bit set), when attempting to execute a MOV on debug registers in V86 mode, a debug exception will be generated instead of the expected general-protection fault.

Workaround: In general, operating systems do not set the GD bit when they are in V86 mode. The GD bit is generally set and used by debuggers. The debug exception handler should check that the exception did not occur in V86 mode before continuing. If the exception did occur in V86 mode, the exception may be directed to the general-protection exception handler.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G28. *Data Breakpoint Exception in a Displacement Relative Near Call May Corrupt EIP*

Problem: If a data breakpoint is programmed at the memory location where the stack push of a near call is performed, the processor will update the stack and ESP appropriately, but may skip the code at the destination of the call. Hence, program execution will continue with the next instruction immediately following the call, instead of the target of the call.

Implication: The failure mechanism for this erratum is that the call would not be taken; therefore, instructions in the called subroutine would not be executed. As a result, any code relying on the execution of the subroutine will behave unpredictably.

Workaround: Do not program a data breakpoint exception on the stack where the push for the near call is performed.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G29. *System Bus ECC Not Functional With 2:1 Ratio*

Problem: If a processor is underclocked at a core frequency to system bus frequency ratio of 2:1 and system bus ECC is enabled, the system bus ECC detection and correction will negatively affect internal timing dependencies.

Implication: If system bus ECC is enabled, and the processor is underclocked at a 2:1 ratio, the system may behave unpredictably due to these timing dependencies.

Workaround: All bus agents that support system bus ECC must disable it when a 2:1 ratio is used.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G30. RDMSR or WRMSR to Invalid MSR Address May Not Cause GP Fault

Problem: The RDMSR and WRMSR instructions allow reading or writing of MSRs (Model Specific Registers) based on the index number placed in ECX. The processor should reject access to any reserved or unimplemented MSRs by generating #GP(0). However, there are some invalid MSR addresses for which the processor will not generate #GP(0).

Implication: For RDMSR, undefined values will be read into EDX:EAX. For WRMSR, undefined processor behavior may result.

Workaround: Do not use invalid MSR addresses with RDMSR or WRMSR.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G31. SYSENTER/SYSEXIT Instructions Can Implicitly Load “Null Segment Selector” to SS and CS Registers

Problem: According to the processor specification, attempting to load a null segment selector into the CS and SS segment registers should generate a General Protection Fault (#GP). Although loading a null segment selector to the other segment registers is allowed, the processor will generate an exception when the segment register holding a null selector is used to access memory.

However, the SYSENTER instruction can implicitly load a null value to the SS segment selector. This can occur if the value in SYSENTER_CS_MSR is between FFF8h and FFFBh when the SYSENTER instruction is executed. This behavior is part of the SYSENTER/SYSEXIT instruction definition; the content of the SYSTEM_CS_MSR is always incremented by 8 before it is loaded into the SS. This operation will set the null bit in the segment selector if a null result is generated, but it does not generate a #GP on the SYSENTER instruction itself. An exception will be generated as expected when the SS register is used to access memory, however.

The SYSEXIT instruction will also exhibit this behavior for both CS and SS when executed with the value in SYSENTER_CS_MSR between FFF0h and FFF3h, or between FFE8h and FFEb, inclusive.

Implication: These instructions are intended for operating system use. If this erratum occurs (and the OS does not ensure that the processor never has a null segment selector in the SS or CS segment registers), the processor's behavior may become unpredictable, possibly resulting in system failure.

Workaround: Do not initialize the SYSTEM_CS_MSR with the values between FFF8h and FFFBh, FFF0h and FFF3h, or FFE8h and FFEb before executing SYSENTER or SYSEXIT.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G32. *PRELOAD Followed by EXTEST Does Not Load Boundary Scan Data*

Problem: According to the IEEE 1149.1 Standard, the EXTEST instruction would use data “typically loaded onto the latched parallel outputs of boundary-scan shift-register stages using the SAMPLE/PRELOAD instruction prior to the selection of the EXTEST instruction.” As a result of this erratum, this method cannot be used to load the data onto the outputs.

Implication: Using the PRELOAD instruction prior to the EXTEST instruction will not produce expected data after the completion of EXTEST.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G33. *Far Jump to New TSS With D-bit Cleared May Cause System Hang*

Problem: A task switch may be performed by executing a far jump through a task gate or to a new Task State Segment (TSS) directly. Normally, when such a jump to a new TSS occurs, the D-bit (which indicates that the page referenced by a Page Table Entry (PTE) has been modified) for the PTE which maps the location of the previous TSS will already be set and the processor will operate as expected. However, if the D-bit is clear at the time of the jump to the new TSS, the processor will hang.

Implication: If an OS is used which can clear the D-bit for system pages, and which jumps to a new TSS on a task switch, then a condition may occur which results in a system hang. Intel has not identified any commercial software which may encounter this condition; this erratum was discovered in a focused testing environment.

Workaround: Ensure that OS code does not clear the D-bit for system pages (including any pages that contain a task gate or TSS). Use task gates rather than jumping to a new TSS when performing a task switch.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G34. *Illegal Opcode During L2 Cache Initialization*

Problem: It is possible for the cache components in the 1-Mbyte and 2-Mbyte Pentium III Xeon processor to power up in a state such that they are not synchronized. During a read under these circumstances, the data in the cache is correct but the processor does not read the data correctly.

Implication: The processor may read invalid data after the cache is enabled during the Power-On Self Test (POST) phase of boot-up, most likely resulting in an invalid opcode being received by the processor, which would generate an invalid opcode exception.

Workaround: Intel recommends that the following BIOS instructions (or equivalent) be added to the Intel L2 Cache initialization module, just prior to enabling the L2 cache via BBL_CR_CTL3 [8]:

```
MOV          ECX, 11Eh          ; MSR (11Eh) is BBL_CR_CTL3
RDMSR                          ; read contents
PUSH         EAX                ; save lower 32 bits
PUSH         EDX                ; save upper 32 bits
AND          AL, 0E1h           ; isolate latency bits
OR           AL, 00Ah           ; set to a desktop latency value
WRMSR                          ; write new value out
POP          EDX                ; restore original value determined
POP          EAX                ; by the BIOS for latency
WRMSR                          ; write it back out
```

IMPORTANT NOTE

The above example code contains stack operations. If the BIOS L2 cache initialization code is executed in a pre-stack environment, the BIOS developer must ensure that the push/pop instruction pairs are replaced with another register save method. Also, the BIOS developer must ensure that the actual BIOS code does not corrupt existing code's register usage.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G35. *Incorrect L2 Cache Line Invalidation*

Problem: In the event of a complex set of internal conditions the processor may invalidate an L2 cache line incorrectly. If this line was previously in the modified state, cache coherency is not maintained. Since timing is critical for these conditions to occur, the following sequence of events MUST occur in order to encounter this erratum:

1. An external snoop occurs.
2. An L1 cache read miss occurs which generates an L2 cache read.
3. This L2 cache read is a miss, causing a modified cache line in the L2 to be scheduled for eviction.
4. An L1 cache write-back (resulting from a separate transaction that happened before event 1) of a different cache line occurs, which causes a second L2 cache miss. The L2 set address is the same as the set address in the first external snoop.
5. Another external snoop occurs.
6. Another L2 cache read occurs; this read may result in either a hit or a miss.
7. The modified cache line in event 3 is written back to main memory, but is not yet invalidated.

The delay between event 3 and event 7 causes the processor to invalidate the wrong L2 cache line.

Implication: Stale data may exist in main memory because the newer, modified data was not written back. Depending on the significance of this modified cache line, results are unpredictable.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G36. *Transmission Error on Cache Read*

Problem: During reads of the L2 cache, the processor may use L2 cache optimizations which result in a data transmission error.

Implication: Data corruption caused by this erratum will result in unpredictable system behavior.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G37. **COMISS/UCOMISS May Not Update EFLAGS Under Certain Conditions**

Problem: COMISS/UCOMISS instructions compare the least significant pairs of packed single-precision floating-point numbers and set the ZF, PF and CF bits in the EFLAGS register accordingly (the OF, SF and AF bits are cleared). The EFLAGS register may not contain the appropriate value for the specified COMISS/UCOMISS operands under the following conditions:

1. The source operand of the COMISS/UCOMISS instruction is from memory that encounters some latency before it is available to the processor

AND

2. An instruction subsequent to the COMISS/UCOMISS is a string instruction (e.g., REP MOVSB) which completes before the UCOMISS/COMISS instruction completes.

Implication: The COMISS/UCOMISS instructions with a register source operand are unaffected by this erratum. The result of the incorrect status of the EFLAGS register may range from no effect to unexpected application or operating system behavior.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G38. **System Hang May Occur With 2:1 Core to Bus Ratio**

Problem: The following conditions must occur for this erratum to be observed:

1. Core frequency to bus frequency ratio must be 2:1.
2. The Chipset must use the most aggressive TRDY assertion allowed by the protocol.
3. Use one of the following instructions: movntq, movntps, or maskmovq.
OR

Do a large number of back to back stores to WC memory.

Implication: The occurrence of this erratum may cause a system hang.

Workaround: Do not use 2:1 core to bus ratio.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G39. **Misaligned Locked Access to APIC Space Results in Hang**

Problem: When the processor's APIC space is accessed with a misaligned locked access a machine check exception is expected. However, the processor's machine check architecture is unable to handle the misaligned locked access.

Implication: If this erratum occurs the processor will hang. Typical usage models for the APIC address space do not use locked accesses. Systems using such a model will not be affected by this erratum.

Workaround: Ensure that all accesses to APIC space are aligned and/or not locked.



Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G40. Potential Loss of Data Coherency During MP Data Ownership Transfer

Problem: In MP systems, processors may be sharing data in different cache lines, referenced as line A and line B in the discussion below. When this erratum occurs (with the following example given for a 2-way MP system with processors noted as 'P0' and 'P1'), P0 contains a shared copy of line B in its L1. P1 has a shared copy of Line A. Each processor must manage the necessary invalidation and snoop cycles before that processor can modify and source the results of any internal writes to the other processor.

There exists a narrow timing window when, if P1 requests a copy of line B it may be supplied by P0 in an Exclusive state which allows P1 to modify the contents of the line with no further external invalidation cycles. In this narrow window P0 may also retire instructions that use the original data present before P1 performed the modification.

Implication: Multiprocessor or threaded application synchronization, required for low-level data sharing, that is implemented via operating system provided synchronization constructs are not affected by this erratum. Applications which rely upon the usage of locked semaphores rather than memory ordering are also unaffected. Uniprocessor systems are not affected by this erratum. The existence of this erratum was discovered during ongoing design reviews but it has not as yet been reproduced in a lab environment. Intel has not identified, to date, any commercially available application or operating system software which is affected by this erratum. If the erratum does occur one processor may execute software with the stale data that was present from the previous shared state rather than the data written more recently by another processor.

Workaround: Deterministic barriers beyond which program variables will not be modified can be achieved via the usage of locked semaphore operations. These should effectively prevent the occurrence of this erratum.

Status: For the processor part numbers affected see the *Pentium® III Processor Identification and Packaging Information* table at the General Information section.

G41. INT 1 Instruction Handler Execution Could Generate a Debug Exception

Problem: If the processor's general detect enable flag is set and an explicit call is made to the interrupt procedure via the INT 1 instruction, the general detect enable flag should be cleared prior to entering the handler. As a result of this erratum, the flag is not cleared prior to entering the handler. If an access is made to the debug registers while inside of the handler, the state of the general detect enable flag will cause a second debug exception to be taken. The second debug exception clears the general detect enable flag and returns control to the handler which is now able to access the debug registers.

Implication: This erratum will generate an unexpected debug exception upon accessing the debug registers while inside of the INT 1 handler.

Workaround: Ignore the second debug exception that is taken as a result of this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G42. *Memory Ordering-Based Synchronization May Cause a Livelock Condition in MP Systems*

Problem: In an MP environment, the following sequence of code (or similar code) in two processors (P0 and P1) may cause them to each enter an infinite loop (livelock condition):

P0		P1
	MOV [xyz], EAX (1)	wait1: MOV EBX, [abc] (2)
	.	CMP EBX, val1 (3)
	.	JNE wait1 (4)
	.	
	MOV [abc], val1 (6)	MOV [abc], val2 (5)
wait0:	MOV EBX, [abc] (7)	
	CMP EBX, val2 (8)	
	JNE wait0 (9)	

NOTE

EAX and EBX can be any general-purpose register. Addresses [abc] and [xyz] can be any location in memory and must be in the same bank of the L1 cache. Variables “val1” and “val2” can be any integer.

The algorithm above involves processors P0 and P1, each of which use loops to keep them synchronized with each other. P1 is looping until instruction (6) in P0 is globally observed. Likewise, P0 will loop until instruction (5) in P1 is globally observed.

The P6 architecture allows for instructions (1) and (7) in P0 to be dispatched to the L1 cache simultaneously. If the two instructions are accessing the same memory bank in the L1 cache, the load (7) will be given higher priority and will complete, blocking instruction (1).

Instructions (8) and (9) may then execute and retire, placing the instruction pointer back to instruction (7). This is due to the condition at the end of the “wait0” loop being false. The livelock scenario can occur if the timing of the wait0 loop execution is such that instruction (7) in P0 is ready for completion every time that instruction (1) tries to complete. Instruction (7) will again have higher priority, preventing the data ([xyz]) in instruction (1) from being written to the L1 cache. This causes instruction (6) in P0 to not complete and the sequence “wait0” to loop infinitely in P0.

A livelock condition also occurs in P1 because instruction (6) in P0 does not complete (blocked by instruction (1) not completing). The problem with this scenario is that P0 should eventually allow for instruction (1) to write its data to the L1 cache. If this occurs, the data in instruction (6) will be written to memory, allowing the conditions in both loops to be true.

Implication: Both processors will be stuck in an infinite loop, leading to a hang condition. Note that if P0 receives any interrupt, the loop timing will be disrupted such that the livelock will be broken. The system timer, a keystroke, or mouse movement can provide an interrupt that will break the livelock.

Workaround: Use a LOCK instruction to force P0 to execute instruction (6) before instruction (7).

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G43. Floating Point Exception Signal May Be Deferred

Problem: A one clock window exists where a pending x87 FP exception that should be signaled on the execution of a CVTTPS2PI, CVTPI2PS, or CVTTPS2PI instruction may be deferred to the next waiting floating point instruction or instruction that would change MMX register state.

Implication: If this erratum occurs the floating point exception will not be handled as expected.

Workaround: Applications that follow Intel programming guidelines (empty all x87 registers before executing MMX technology instructions) will not be affected by this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G44. System Bus Address Parity Checking May Report False AERR#

Problem: The processor's address parity error detection circuit may fail to meet its frequency timing specification under certain environmental conditions. At the high end of the temperature specification and/or the low end of the voltage range, the processor may report false address parity errors.

Implication: If the system has AERR# drive enabled (bit [3] of the EBL_CR_POWERON register set to '1') spurious address detection and reporting may occur. In some system configurations BINIT# may be asserted on the system bus. This may cause some systems to generate a machine check exception and in others may cause a reboot.

Workaround: Disable AERR# drive from the processor. AERR# drive may be disabled by clearing bit [3] in the EBL_CR_POWERON register. In addition, if the chipset allows, AERR# drive should be enabled from the chipset and AERR# observation enabled on the processor. AERR# observation on the processor is enabled by asserting A8# on the active-to-inactive transition of RESET#.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G45. Processor May Assert DRDY# on a Write with No Data

Problem: When a MASKMOVQ instruction is misaligned across a chunk boundary in a way that one chunk has a mask of all 0's, the processor will initiate two partial write transactions with one having all byte enables deasserted. Under these conditions, the expected behavior of the processor would be to perform both write transactions, but to deassert DRDY# during the transaction which has no byte enables asserted. As a result of this erratum, DRDY# is asserted even though no data is being transferred.

Implication: The implications of this erratum depend on the bus agents ability to handle this erroneous DRDY# assertion. If a bus agent cannot handle a DRDY# assertion in this situation, or attempts to use the invalid data on the bus during this transaction, unpredictable system behavior could result.

Workaround: A system which can accept a DRDY# assertion during a write with no data will not be affected by this erratum. In addition, this erratum will not occur if the MASKMOVQ is aligned.

Status: For the steppings affected see the *Summary Table of Changes* at the beginning of this section.

G46. Thermal Sensor Leakage Current May Exceed Specification

Problem: The thermal sensor in the Pentium III Xeon processor cartridge violates the maximum input leakage current specification in Table 9 of the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet (10 μ A).

Implication: The thermal sensor incorporates input protection diodes on the SMBCLK and SMBDAT signals for ESD protection. The protection diodes can potentially clamp these lines to ~ 0.6 V when the $V_{CCSMBus}$ voltage supply to the cartridge is powered off. Hence, when $V_{CCSMBus}$ is powered down, the Pentium III Xeon processor thermal sensor may prevent SMBus transactions originating from SMBus devices powered by a different voltage source, but on the same SMBus, from occurring. SMBus devices that are powered from different voltage power planes are not typically located on the same SMBus. Designs using this isolation technique will not be affected by this erratum.

Workaround: If it is desired to have SMBus devices powered by a source other than $V_{CCSMBus}$, these devices must be isolated from the Pentium III Xeon processor SMBus to ensure that this erratum does not occur.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G47. GP# Fault on WRMSR to ROB_CR_BKUPTMPDR6

Problem: Writing a '1' to unimplemented bit(s) in the ROB_CR_BKUPTMPDR6 MSR (offset 1E0h) will result in a general protection fault (GP#).

Implication: The normal process used to write an MSR is to read the MSR using RDMSR, modify the bit(s) of interest, and then to write the MSR using WRMSR. Because of this erratum, this process may result in a GP# fault when used to modify the ROB_CR_BKUPTMPDR6 MSR.

Workaround: When writing to ROB_CR_BKUPTMPDR6 all unimplemented bits must be '0.' Implemented bits may be set as '0' or '1' as desired.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G48. Heavy L2 Cache Traffic Results in Noise on the TCK Signal

Problem: Pentium III Xeon processors under heavy L2 cache traffic may induce noise on the TCK signal. If the system is also accessing the processor through the Test Access Port, erroneous operation may occur.

Implication: Access to the Test Access Port may fail if access is simultaneous with heavy L2 cache traffic.

Workaround: Ensure that any access to the Test Access Port does not occur simultaneously with heavy L2 cache traffic.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G49. *Machine Check Exception May Occur Due to Improper Line Eviction In The IFU*

Problem: The Pentium III Xeon processor is designed to signal an unrecoverable Machine Check Exception (MCE) as a consistency checking mechanism. Under a complex set of circumstances involving multiple speculative branches and memory accesses there exists a one-cycle long window in which the processor may signal a MCE in the Instruction Fetch Unit (IFU) because instructions previously decoded have been evicted from the IFU. The one cycle long window is opened when an opportunistic fetch receives a partial hit on a previously executed but not as yet completed store resident in the store buffer. The resulting partial hit erroneously causes the eviction of a line from the IFU at a time when the processor is expecting the line to still be present. If the MCE for this particular IFU event is disabled, execution will continue normally.

Implication: While this erratum may occur on a system with any number of Pentium III Xeon processors, the probability of occurrence increases with the number of processors. If this erratum does occur, a machine check exception will result. Note systems that implement an operating system that does not enable the Machine Check Architecture will be completely unaffected by this erratum (e.g., Windows* 95 and Windows 98).

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



G50. Machine Check Exception May Occur During L2 Cache Initialization

Problem: It is possible for a multiple bit ECC error to be induced by the random state of the TAG bits in the processor's cache component(s) at power up. If this multiple bit ECC error occurs, it may, depending on system configuration, result in a machine check exception error occurring during cache initialization.

Implication: This erratum may occur on a system with Pentium III Xeon processors which utilize TAG ECC within the cache component(s). The occurrence of this erratum is dependent on the specific cache initialization sequence. If this erratum occurs, a machine check exception will result.

Workaround: Executing the following code sequence in BIOS prior to enabling the processor's L1 cache will ensure that this erratum does not occur. In the case of a multiprocessor system, ensure that the BIOS workaround code executes on each processor individually prior to enabling the L1 cache:

```

;-----;
; PUSHAD ; Save registers.
;-----;
; Get feature bits now and check for processor type (family/model).
;-----;
MOV     EAX, 1h ; CPUID version/feature parameter.
CPUID   ; Execute CPUID instruction.
CMP     AX, 672h ; Check for Pentium III Xeon processor, Model 7,
JL      SkipWorkaround ; If not B0 stepping or greater of this processor..
CMP     AX, 680h ; ..exit, skipping the workaround.
JGE     SkipWorkaround
;-----;
; Check Platform ID bits now and only execute code if Slot 2.
;-----;
MOV     ECX, BBL_CR_OVRD ; MSR 17h
RDMSR   ;
AND     EDX, 001C0000h ; check Platform ID bits to ensure that ..
CMP     EDX, 00080000h ; .. it is the SC330 version of the CPUID..
JNE     SkipWorkaround ; .. If not, exit
;-----;
; Ensure that L2 cache is disabled. The L2 cache is disabled by clearing BBL_CR_CTL3 bit 8.
; Also, we write BBL_CR_CTL3 controlled L2 features to default WITHOUT disturbing
; RESERVED or read only register bits. The previous BBL_CR_CTL3 settings are restored
; on completion.
;-----;
MOV     ECX, BBL_CR_CTL3 ; MSR 11Eh
RDMSR   ;
PUSH    EAX ; Save original value.
PUSH    EDX
AND     EAX, 0FF881E00h ; Clear address space, disable CSER checking, clear cache
; size, disable L2, disable CRTN parity checking, disable
; victim address parity checking, disable ECC, L2 not
; configured, and clear latency setting
OR      EAX, 00001001Fh ; Set 2MB, latency, and L2 configured
WRMSR   ; Flush L2 tags.
INVD    ; Restore original value
POP     EDX
POP     EAX
WRMSR   ;
SkipWorkaround:
;-----;
POPAD ; Restore registers.
;-----;
```

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G51. Snoop Request May Cause DBSY# Hang

Problem: A small window of time exists in which a snoop request originating from a bus agent to a processor with one or more outstanding memory transactions may cause the processor to assert DBSY# without issuing a corresponding bus transaction, causing the processor to hang (livelock). The exact circumstances are complex, and include the relative timing of internal processor functions with the snoop request from a bus agent.

Implication: This erratum may occur on a system with any number of processors. However, the probability of occurrence increases with the number of processors. If this erratum does occur, the system will hang with DBSY# asserted. At this point, the system requires a hard reset.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G52. Performance Counters Include Streaming SIMD Extensions L1 Prefetch

Problem: The processor allows the measurement of the frequency and duration of numerous different internal and bus related events (see the *Intel Architecture Software Developer's Manual, Volume 3*, for more details). The Streaming SIMD Extension (SSE) architecture provides a mechanism to pre-load data into the L1 cache, bypassing the L2 cache. The number of these L1 pre-loads measured by the performance monitoring logic will incorrectly be included in the count of "L2_LINES_IN" (24H) events and "L2_LINES_OUT" (26H) events.

Implication: If application software is run which utilizes the SSE L1 prefetch feature, the count of "L2_LINES_IN" (24H) and "L2_LINES_OUT" (26H) will read a value that is greater than the correct value.

Workaround: The correct value of "L2_LINES_IN" and "L2_LINES_OUT" may be calculated by subtracting the value of the "MMX_PRE_MISS" (4BH) from each of these registers.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G53. Lower Bits of SMRAM SMBASE Register Cannot Be Written With an ITP

Problem: The System Management Base (SMBASE) register (7EF8H) stores the starting address of the System Management RAM (SMRAM). This register is used by the processor when it is in System Management Mode (SMM), and its contents serve as the memory base for code execution and data storage. The 32-bit SMBASE register can normally be programmed to any value. When programmed with an In-Target Probe (ITP), however, any attempt to set the lower 11 bits of SMBASE to anything other than zeros via the WRMSR instruction will cause the attempted write to fail.

Implication: When set via an ITP, any attempt to relocate SMRAM space must be made with 2-Kbyte alignment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G54. Task Switch Caused by Page Fault May Cause Wrong PTE and PDE Access Bit to be Set

Problem: If an operating system services a page fault through a Task Gate, the access bit may be set for an incorrect page table/directory entry.

Implication: An operating system which uses a Task Gate for its page fault handler may encounter this erratum. The effect of the erratum depends on the alignment of the Task State Segment (TSS), and ranges from no anomalous behavior to unexpected errors. Intel is not aware of any commercial operating systems which use a Task Gate to handle page faults. Task gates used for other purposes (NMI, Machine Check, or Double Fault) do not cause this erratum.

Workaround: The operating system may alternately use an Interrupt Gate or a Call Gate rather than a Task Gate.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G55. Cross-Modifying Code Operations on a Jump Instruction May Cause General Protection Fault

Problem: The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called Cross-Modifying Code (XMC). Software using XMC to modify the offset of an execution transfer instruction i.e. Jump, Call etc., without a synchronizing instruction may cause a General Protection Fault (GPF) when the offset splits a cache line boundary.

Implication: Any application creating a GPF would be terminated by the operating system.

Workaround: Programmers should use the cross-modifying code synchronization algorithm as detailed in Volume 3 of the *Intel Architecture Software Developer's Manual*, Section 7.1.3, in order to avoid this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G56. V_{IH} Specification Deviation on Processor RS0# Input


Problem: The RS0# signal of the 550 MHz Pentium III Xeon processor with 512K and 1M of cache has been found to deviate from the published specification for the minimum input high voltage, $V_{IH(MIN)}$, as published in the *Pentium® III Xeon™ Processor at 500 MHz and 550 MHz Datasheet*.

The published AGTL+ AC specification for $V_{IH(MIN)}$ is $(2/3 \cdot V_{TT}) + 0.1$ Volts.

Intel has characterized the new AGTL+ AC specification for $V_{IH(MIN)}$ on the processor's RS0# signal to be: $(2/3 \cdot V_{TT}) + 0.22$ Volts within 500 picoseconds following the rising edge of BCLK crossing the 1.25-Volt threshold. See below.



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE



Implication: The deviation from the published specification results in decreased noise margin on the RS0# processor input. This decreased noise margin has been observed to be particularly detrimental in system environments where the RS0# processor input is being driven by an I/O buffer with a pull-up device (such as a p-channel transistor) that is switched off mid-cycle. Systems which employ I/O buffers that keep the pull-up device on for the entire cycle are significantly less affected by this erratum. The reduction in noise margin may lead to the generation of BINIT# (when enabled) resulting in a system reset, otherwise this condition may lead to a system "hang" condition.

Workaround: If a system design cannot meet the specification deviation, the following areas should be carefully examined to assure that additional margin exists in order to compensate for any deficit observed with respect to the newly characterized $V_{IH(MIN)}$ specification:

1. Consider increasing V_{TT} to a higher voltage than 1.5 Volts. Higher V_{TT} leads to improved noise margin by increasing the difference between V_{TT} and V_{REF} . For each 10-mV increase in V_{TT} , an additional 3.3-mV improvement in margin is achieved. Care should be exercised so as to avoid exceeding the maximum V_{TT} specification of 1.635 Volts.
2. For crow's foot front side bus topology, reduce the value of the termination resistor attached to the RS0# network at the main trunk from the original recommended value of 150Ω to 50Ω . If no resistor exists at this junction, a $50\text{-}\Omega$ resistor may be added at this point. Other topologies should be examined to determine if reduction in the termination impedance at the source achieves additional margin to the newly characterized $V_{IH(MIN)}$ specification.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G1AP. ***APIC Access to Cacheable Memory Causes Shutdown***

Problem: APIC operations which access memory with any type other than uncacheable (UC) are illegal. If an APIC operation to a memory type other than UC occurs and Machine Check Exceptions (MCEs) are disabled, the processor will enter shutdown after such an access. If MCEs are enabled, an MCE will occur. However, in this circumstance, a second MCE will be signaled. The second MCE signal will cause the Pentium III Xeon processor to enter shutdown.

Implication: Recovery from a PIC access to cacheable memory will not be successful. Software that accesses only UC type memory during APIC operations will not encounter this erratum.

Workaround: Ensure that the memory space to which PIC accesses can be made is marked as type UC (uncacheable) in the memory type range registers (MTRRs) to avoid this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G2AP. ***MP Systems May Hang Due to Catastrophic Errors During BSP Determination***

Problem: In MP systems, a catastrophic error during the bootstrap processor (BSP) determination process should cause the assertion of IERR#. If the catastrophic error is due to the APIC data bus being stuck at electrical zero, then the system hangs without asserting IERR#.

Implication: MP systems may hang during boot due to a catastrophic error. This erratum has not been observed to date in a typical commercial system, but was found during focused system testing using a grounded APIC data bus.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

G3AP. ***Write to Mask LVT (Programmed as EXTINT) Will Not Deassert Outstanding Interrupt***

Problem: If the APIC subsystem is configured in Virtual Wire Mode implemented through the local APIC, (i.e., the 8259 INTR signal is connected to LINT0 and LVT1's interrupt delivery mode field is programmed as EXTINT), a write to LVT1 intended to mask interrupts will not deassert the internal interrupt source if the external LINT0 signal is already asserted. The interrupt will be erroneously posted to the Pentium III Xeon processor despite the attempt to mask it via the LVT.

Implication: Because of the masking attempt, interrupts may be generated when the system software expects no interrupts to be posted.

Workaround: Software can issue a write to the 8259A interrupt mask register to deassert the LINT0 interrupt level, followed by a read to the controller to ensure that the LINT0 signal has been deasserted. Once this is ensured, software may then issue the write to mask LVT entry 1.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



DOCUMENTATION CHANGES

The Documentation Changes listed in this section applies to the following documents:

- *P6 Family of Processors Hardware Developer's Manual*
- *Pentium® III Xeon™ Processor at 500 and 550 MHz datasheet*
- *Pentium® III Xeon™ Processor at 600 MHz, 667 MHz and 733 MHz datasheet*
- *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3*

All Documentation Changes will be incorporated into a future version of the appropriate Pentium III Xeon processor documentation.

G1. STPCLK# Pin Definition

The *P6 Family of Processors Hardware Developer's Manual* and the *Pentium® III Xeon™ Processor at 500 and 550 MHz datasheet* both have incorrect definitions of the STPCLK# pin in their alphabetical signal listing. These documents incorrectly state:

The processor continues to snoop bus transactions and *service interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units and resumes execution.

They should state:

The processor continues to snoop bus transactions and *may latch interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units, resumes execution, *and services any pending interrupts*.

G2. Invalidating Caches and TLBs

Section 2.6.4 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, incorrectly states:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and TLBs and sends a signal to the external caches indicating that they should be invalidated also.

It should state:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and sends a signal to the external caches indicating that they should be invalidated also.

G3. Handling of Self-Modifying and Cross-Modifying Code

Section 7.1.3 paragraph 1. of the *Intel Architecture Software Developer's Manual Vol. 3: System Programming* incorrectly states:

The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **should** be chosen.

It should state:

The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **must** be chosen.

Section 7.1.3 paragraph 6 of the *Intel Architecture Software Developer's Manual Vol. 3: System Programming* incorrectly states:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **should** be implemented.

It should state:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **must** be implemented.



SPECIFICATION CLARIFICATIONS

The Specification Clarifications listed in this section applies to the following documents:

- *P6 Family of Processors Hardware Developer's Manual*
- *Pentium® III Xeon™ Processor at 500 and 550 MHz datasheet*
- *Pentium® III Xeon™ Processor at 600 MHz, 667 MHz and 733 MHz datasheet*
- *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3*

All Specification Clarifications will be incorporated into a future version of the appropriate Pentium III Xeon processor documentation.

G1. MTRR Initialization Clarification

The following sentence should be added to the end of the first paragraph of Section 9.11.5 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*: "The MTRRs must be disabled prior to initialization or modification."

G2. Floating Point Opcode Clarification

Section 3.2 of the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference*, provides detailed descriptions of each Intel Architecture instruction. For some instructions, the clarification phrase below needs to be either added to their existing "Comments" section or a "Comments" section needs to be created with the clarification phrase. The phrase is as follows:

The (**Instruction** shown in the center column of the table below) instruction is actually a combination of two instructions – the wait instruction followed by (**Instruction** shown in the table). If the (**Instruction** shown in the table) instruction should fault in some way (e.g., page fault), the value of EIP that is passed to the fault handler will be equal to the EIP of the first instruction plus one (i.e., the EIP of the second of the pair of instructions). The FWAIT portion of the combined instruction will have completed execution and will typically not be, nor need to be, re-executed after the fault handler is completed.



PENTIUM® III XEON™ PROCESSOR SPECIFICATION UPDATE

The following table lists the affected instructions and the location of the clarification phrase:

Instruction Set Reference Section	Opcode	Instruction	Clarification	Page
FCLEX/FNCLEX-Clear Expectation	9B DB E2	FCLEX	Add “Comments” section with clarification phrase	3-177
FINIT/FNINIT-Initialize Floating-Point Unit	9B DB E3	FINIT	Add clarification phrase to existing “Comments” section	3-204
FSAVE/FNSAVE-Store FPU State	9B DD /6	FSAVE m94/108byte	Add clarification phrase to existing “Comments” section	3-237
FSTCW/FNSTCW-Store Control Word	9B D9 /7	FSTCW m2byte	Add “Comments” section with clarification phrase	3-250
FSTENV/FNSTENV- Store FPU Environment	9B D9 /6	FSTENV m14/28byte	Add “Comments” section with clarification phrase	3-253
FSTSW/FNSTSW-Store Status Word	9B DD /7	FSTSW m2byte	Add “Comments” section with clarification phrase	3-256
	9B DF E0	FSTSW AX		



SPECIFICATION CHANGES

The Specification Changes listed in this section applies to the following documents:

- *P6 Family of Processors Hardware Developer's Manual*
- *Pentium® III Xeon™ Processor at 500 and 550 MHz datasheet*
- *Pentium® III Xeon™ Processor at 600 MHz, 667 MHz and 733 MHz datasheet*
- *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3*

All Specification Changes will be incorporated into a future version of the appropriate Pentium III Xeon processor documentation.

The latest datasheet was published January 1998.

G1. Locks Across Cache Line Boundary Disable Bit Added

In the Pentium III Xeon processor, setting bit 31 of the Model Specific Register (MSR) at address 33h to '1' will prevent LOCK# from being asserted when locked transactions which are split across a cache line boundary are issued from the processor. This bit is disabled by default. In the default state ('0'), unaligned data issued in a locked sequence by the processor will have atomicity with the LOCK# signal asserted. When the bit is set, any transactions issued which split a cache line boundary will not have the LOCK# signal asserted, and no atomicity can be guaranteed between the reads and writes in the sequence. Locked sequences which do not split a cache line boundary will still follow the normal LOCK# protocol with this bit set.

Pentium® III Xeon™ Processor Stepping	500-MHz B0	550-MHz C0
500-MHz B0	1	X
550-MHz C0	X	1

G2. RESET# Pin Definition

The *P6 Family of Processors Hardware Developer's Manual* and the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet both have incorrect definitions of the RESET# pin in their alphabetical signal listing. These documents incorrectly state:

RESET# must remain active for one microsecond for a "warm" Reset; for a Power-on Reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.

They should state:

For a Power-on or "warm" reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.



G3. RESET# Pin Definition

The *P6 Family of Processors Hardware Developer's Manual* and the *Pentium® III Processor at 450 MHz and 500 MHz* datasheet both have incorrect definitions of the RESET# pin in their alphabetical signal listing. These documents incorrectly state:

RESET# must remain active for one microsecond for a "warm" Reset; for a Power-on Reset, RESET# must stay active for at least one millisecond after V_{CCCORE} and CLK have reached their proper specifications.

They should state:

For a Power-on or "warm" reset, RESET# must stay active for at least one millisecond after V_{CCCORE} and CLK have reached their proper specifications.



UNITED STATES, Intel Corporation
2200 Mission College Blvd., P.O. Box 58119, Santa Clara, CA 95052-8119
Tel: +1 408 765-8080

JAPAN, Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi, Ibaraki-ken 300-26
Tel: + 81-29847-8522

FRANCE, Intel Corporation S.A.R.L.
1, Quai de Grenelle, 75015 Paris
Tel: +33 1-45717171

UNITED KINGDOM, Intel Corporation (U.K.) Ltd.
Pipers Way, Swindon, Wiltshire, England SN3 1RJ
Tel: +44 1-793-641440

GERMANY, Intel GmbH
Dornacher Strasse 1
85622 Feldkirchen/ Muenchen
Tel: +49 89/99143-0

HONG KONG, Intel Semiconductor Ltd.
32/F Two Pacific Place, 88 Queensway, Central
Tel: +852 2844-4555

CANADA, Intel Semiconductor of Canada, Ltd.
190 Attwell Drive, Suite 500
Rexdale, Ontario M9W 6H8
Tel: +416 675-2438

BRAZIL, Intel Semicondutores do Brasil
Centro Empresarial Nações Unidas - Edifício Torre Oeste
Av. das Nações Unidas, 12.901 - 18o. andar - Brooklin Novo
04578.000 São Paulo - S.P. - Brasil
Tel: +55-11-5505-2296